Use of Concave Corners in the Segmentation of Embryological Datasets

Kieran Rafferty [1], Sarah Drury [2], Geraldine Hartshorne [2], Silvester Czanner [*1]

[*1]School of Computing, Mathematics and Digital Technology, Manchester Metropolitan University, Manchester , M1 5GD, UK

[2]Division of Reproductive Health, Clinical Sciences Research Laboratories, Warwick Medical School, University of Warwick, Coventry, CV2 2DX, UK

[1]Kieran.Rafferty@stu.mmu.ac.uk; [2]S.L.Drury@warwick.ac.uk; [2]Geraldine.Hartshorne@warwick.ac.uk;

[*1] S.Czanner@mmu.ac.uk

Abstract

Embryologists require identifiable factors when analysing embryos grown in-vitro that can better predict an embryos successful development in the uterus. It has been proposed that the volume of various embryo features, such as blastomeres, within a developing of embryo could be used as an indicator of viability. An initial interface has been developed using existing image processing tools to segment z-stacks produced via confocal microscopy and create a set of volumes for analysis. A further addition to these tools was developed with the intent of improving the accuracy of the segmentation process; specifically, in identifying individual blastomeres whose proximity to neighbouring blastomeres resulted in the false identification of a single segment. A method has been proposed that uses a contour produced by an earlier segmentation procedure and identifies concave regions along the contour. These are interpreted as probable indicators of the location of boundaries between different blastomeres. By using points along the contour in these concave regions we attempt to extrapolate the boundaries between blastomeres.

Keywords

Embryological datasets; Image segmentation; Embryology; Blastomeres; Concave corners

Introduction

Assisted reproductive technology (ART) is a term that covers a range of methods used primarily in infertility treatments achieving pregnancy through artificial or partially artificial means. Methods of ART include fertility medication, which is used to stimulate follicle development of the ovary and in vitro fertilisation (IVF), where an egg is fertilised outside of the body in a laboratory setting *in vitro* and then transferred to the patients uterus. The emphasis of current research in the field of IVF is to determine the qualities by which an embryo can be evaluated that would indicate it is likely to result in a successful pregnancy. (Filho, 2010) refers to a study that compared the grades compiled from a number of practicing embryologists against a control. Substantial differences between morphological scores were observed by as much as two grades, despite using the same grading system. The limited time between insemination and implantation is a large determining factor in the ability of embryologists to make quantified assessments of developing embryos. Embryos at the blastocyst stage may have greater than a hundred constituent cells and accurately measuring the size of each would be a long and tedious process. The development of computer algorithms to take these measurements would be ideal in eliminating this hurdle. Using image processing it is hoped that potential indicators can be quantified and accurately measured.

Grading of the blastocyst has become more popular in recent years (Filho, 2010). The blastocyst stage of embryo development is usually reached around five days after insemination and some nonviable embryos may undergo apoptosis, automatically eliminating themselves from the embryo selection process. By waiting for embryos to reach the blastocyst stage of development embryologists have a better idea of the potential viability of an embryo. Despite this, embryologists still require further identifiable factors that can better predict an embryos successful development in the uterus.

A large number of studies that have attempted to automatically segment embryos have used images produced from HMC optical to do so (Morales et al 2008, Pederson et al 2003, Giusti et al 2010, Karlsson et al 2004). Fewer have attempted to use confocal imaging though further examples of confocal microscopy segmentation are found outside of

embryology (Smochina et al 2011).

Morales et al (2008) investigated techniques to measure Zona Pellucida thickness variation using an active contour processing filter deciding the ability for snake segmentation to ignore small artefacts and adapt to the shape of the Zona Pellucida regardless of whether this was elliptical or circular.

Pederson et al (2003) developed a technique for modelling the relative position of blastomeres and their 3D shape within the embryo at the second day after fertilisation. Again, using HMC they detected the contours of blastomeres and outlining a method to overcome the issue that the contour of blastomeres will may be obscured by other blastomeres at a given depth. After presenting visually positive resulting images, they continue to explain how these support the existing theories of embryo structure without explaining how this has been quantifiably supported.

Tassy et al (2006) discussed the development of a program which can be used to analyse the 3D models produced as a result of segmentation of confocal laser microscopy image stacks. The actual segmentation is only mentioned as an afterthought where manual reconstruction of a 44 cell embryo is said to potentially take up to a full day.

Smochina et al (2011) discuss reasons why confocal microscopy is difficult to segment. These are listed as differing grey values for the background across different regions of the image; Low contrast, weak boundaries of out of focus objects; variation of objects shape/size and orientation; finally chained or touching objects having indistinguishable boundaries.

Methodology

A number of preliminary tests were run on a variety of region growing and edge detection filters. Due to the level of noise in many of the images, it became apparent that using edge detection filters as a method of segmentation was infeasible. Images were over-segmented with edges rarely corresponding in an identifiable pattern with cell perimeters. Because the intent is to produce outlines of objects that can be transformed into vector lines and subsequently polygons, active contours were also considered. Active Contours operated on multiple iterations which increased the processing time of images and impacting the user's feedback when changes were made. Additionally, many of the Active Contour filters expected a template shape to be provided and due to

the potential variance between cell shapes (due to either their orientation in a given slice or distortion of the shape by neighbouring cells) it was hoped that this requirement could be avoided by relying on region growing segmentation.

Two region growing segmentation filters were employed. The first employed Neighbourhood Connected Thresholding. Here, two thresholds were defined, a minimum and a maximum. A pixel was included in a region if all of its neighbours fell within that range. This filter compensated for noise fluctuations because small individual points with significantly different intensities would be averaged out.

The second filter used Confidence Connected Thresholding. With this filter, a point or 'seed' on the image could be specified and the mean and standard deviation of pixel intensity values in a specified neighbourhood around the seed was calculated. The result of multiplying the standard deviation by a user-specified value means a range around the mean is defined. Neighbouring pixels, whose intensity values fall inside the range, can be accepted into the region. When all the pixels that satisfy this criterion have been included, the first iteration of the filter is considered complete. Subsequent iterations recalculate the mean and standard deviation using all the pixels currently included in the region. The pixels neighbouring those currently in the region are then evaluated in the same was as before to determine their inclusion in the region.

$$I(P) \in R \longrightarrow P(i) = [m - f\sigma, m + f\sigma]$$

I is the image, P is a pixel, R is a region of the image, i is the intensity of a pixel, m is the mean, σ is the standard deviation, and f is a user supplied value. Increasing f means a much wider range of values that a pixel's intensity can lie within to be included in a given region.

The Neighbourhood Connected Threshold operates globally on the image. In order to minimise the variability of segment size between slices, the use of a global threshold was considered useful. By doing so, a consistent threshold can be established and applied to each slice with only minor adjustments to compensate for photo-bleaching as we move through the slices. When over-segmentation occurs in a localised region, we can switch to the confidence connected threshold filter in order to control the growth of the intended region.

The edge detection performs a fairly limited role in the process. As mentioned earlier, the amount of noise present in the images interfered with the edge detectors ability to produce complete, cell-shaped contours. However, the output of the region growing filter was a binary colour output of segmented regions. In order to produce outlines of the various shapes that could be transformed into vector points, and subsequently polygon shapes, it was necessary to apply an edge detection filter to the region growing output (FIG 1).

FIG. 1 THE IMAGE IN THE LEFT SLIDE IS A BLASTOMERE IN THE UNALTERED IMAGE. THE CENTRE IMAGE IS THE RESULT OF A CONFIDENCE CONNECTED THRESHOLD FILTER. THE RIGHT IMAGE THE RESULT OF

Concave Corner Detection

Your The primary issue that became quickly apparent was the difficulty in identifying individual cells when they were closely grouped. Both region growing algorithms essentially calculated if a pixel fell within a given range. Cells focused in the same slice display little variance in their overall intensity. The dye used in confocal microscopy adheres to various parts of the cell differently. Nuclei, for example, may be identifiable within a cell because it will display a different intensity to the cytoplasm. The noise that is attributed to the image acquisition process also affects the overall intensity of a cell. The smoothing filter helps to compensate for the intra-cell variance, and the region growing filters also accept pixels lying between a given threshold range. Because the variance of intensity between cells is small (each cell is subject to noise and each cell contains nuclei) this is sufficient to identify multiple cells from the background simultaneously. It is often possible for this to also identify in-focus cell shapes (higher intensity pixels) from out of focus (lower intensity pixels). However, the nature of the region growing filters makes it very difficult for them to distinguish neighbouring cells from one another and will often include them in the same segment.

Consider the confidence connected threshold filter which grows the region from a specified seed. In order to compensate for the intra-cell variance the multiplier is increased in order to include a wider range of

intensities than those found at the starting pixel. If two cells are immediately neighbouring one another, the edge that exists between them will have a more subtle gradient than the edge that exists between a cell and the background. In fact, it is likely that the change in intensity that identifies the edge lies beneath the threshold that was established to compensate for the intra-cell variance. Rather than being segmented as two separate shapes, a single shape is identified. Reducing the multiplier results in over-segmentation as the intra-cell variance is no longer compensated for. This problem is present at every stage of embryo development. Grading systems of cleavage stage embryos often revolve around the proximity of blastomeres (sources/references), while the trophectoderm and inner cell mass of blastocysts are both characterised by numerous adjacent cells.
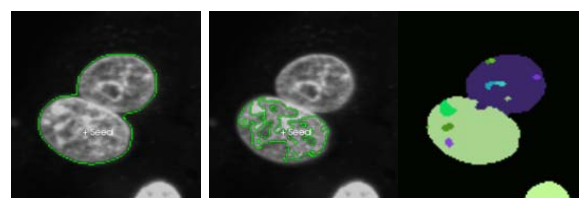

FIG. 2A) TWO NEIGHBOURING CELLS IN CONTACT WITH ONE ANOTHER AND THE RESULT OF CONFIDENCE CONNECTED THRESHOLDING. SEED MARKER IDENTIFIES CELL INTENDED FOR SEGMENTATION. A SUFFICIENTLY HIGH MULTIPLIER TO COMPENSATE FOR THE INTRA-CELL INTENSITY VARIATION RESULTS IN THE CELL BEING GROUPED IN THE SAME REGION WITH AN ADJACENT CELL. FIG 2B) REDUCING THIS MULTIPLIER TO EXCLUDE THIS NEIGHBOURING CELL RESULTS IN OVER SEGMENTATION OF THE INTENDED CELL. FIG 2C) RESULT OF WATER SHED SEGMENTATION

Initially this problem was overcome by using a Watershed Segmentation filter. Although often successful, the filter was extremely processor intensive and best results were observed after more than forty iterations. Rendering a new image after changing a variable could take up to a minute on a slow machine. Again over-segmentation remained a problem though not nearly as severe and greater control appeared to be available. Successes varied from case to case (compare FIG 2c and FIG 3b).
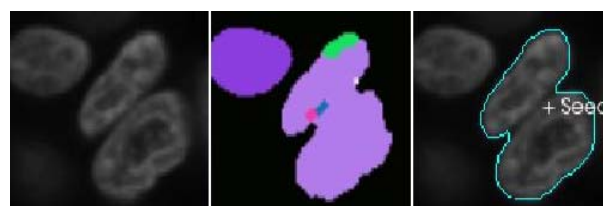

FIG. 3A), 3B), 3C) FAILED ATTEMPTS OF BOTH WATERSHED AND CONFIDENCE CONNECTED THRESHOLDING TO DISTINGUISH TWO NEIGHBOURING CELLS

Implementation of Concave Corners

To overcome this issue, a post-segmentation process was decided upon. Whereas an active contour may require a suggested template to which the algorithm should aim to grow towards, the intention here was estimate the shape that should be aimed towards from the available information. When two cells are in close enough proximity to one another for the region growing filter to detect them as a single segment, the contour that is drawn contains two corners either side of the contact boundary. By identifying these corners, it would be possible to identify segments that are likely to be two or more intended segments falsely identified as one. Furthermore, as cells come into contact with one another their shape deforms. It seems possible that by knowing the location of the corners, in combination with the shape of the cell where the cells meet (i.e. the shape of the corners), we would be able to predict the edge that defines the two cells.

The process was initially divided into the following stages:

1. Trace a contour and identify all points which form a concave angle with adjacent points

2. Eliminate any set of points beneath a given size that may be considered insignificant fluctuations in an object's contour

3. Extrapolate dividing edge from available points in the corner.

Identifying what constituted a corner presented two problems. First, the angle between two lines was calculated via the law of cosines. This produced an answer in the range of 0-180. It was necessary to use ray casting to determine whether this angle was on the inside or the outside of the shape. The second, more significant problem was that the contour of cells was imperfect. That is, they were not a consistent curve as you would expect on a circle or eclipse, but contained small bends that could be falsely interpreted as corners. This was often a result of the maximum pixel resolution of the source image, or noise that had been incorrectly drawn around. It was therefore necessary to distinguish what was a fluctuation and what was part of a corner. Requiring a given point to meet a certain angle was considered but an actual corner may be very obtuse and fail to meet this threshold. The best resolution appeared to be requiring that the corner contained a minimum number of concave angles. Here we will define a concave angle as where three points A, B and C produce an angle on the exterior of the

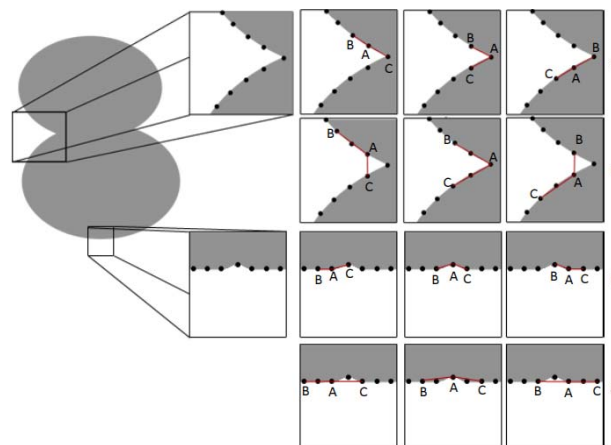polygon shape at point A where line AB and AC meet, and which is less than 180 degrees.

FIG 5 MINIMUM CORNER SIZE REQUIRED TO IGNORE FLUCTUATIONS. MINIMUM DISTANCE BETWEEN POINTS REQUIRED TO MAKE SURE THE NUMBER OF POINTS IN A CORNER IS > 1

Obviously, the minimum number of concave angles would need to be dynamic because smaller cells would have fewer points in their contour. The problem that presented itself with this solution was that stepping over each point and measuring the angle between itself and the points either side of it would often only produce a single concave angle. Consider the corner presented in FIG 5. Each black dot is a point along the shape's contour. In the top row, three points have been labelled A, B, and C. A is the point we are currently interested in. An angle is formed at A between connecting lines BA and CA (shown in red). When points B and C are one point away from A, only one angle in the top row is determined to be concave (second from right, top row). In the third row, where there is a small fluctuation in the contour and the points B and C are a single step away from A, we find as many concave angles as we did in the first row despite the two regions being clearly distinct.

In the second row points B and C are two points away from A. This results in a greater number of concave angles being identified, perhaps enough to signify we have located a corner that is of interest to us. Advantageously, the same increase in the distance of points B and C from A still results in only one concave corner in row four as opposed to three by the scenario displayed in row two. By increasing the distance between points, larger fluctuations in a contour are ignored as lines AB and AC fail to create a concave angle or fail to create a sufficient number of consecutive points with concave angles. Of course, increasing the distance over a certain amount would

cause points B and C to wrap to parts of the shape away from the region of interest, perhaps on the opposite side of the shape, and form angles that falsely represent the overall contour. Fluctuations within a corner that may interrupt a series of points and result in a failure to meet the threshold are also largely avoided, changing the angle but not severely enough to affect successful identification. The advantage of this was that significantly bigger concave regions could be ignored, often until only those of interest remained. Like the threshold that determines the minimum number of points that must make up a corner, the distance that points are from one another must have a maximum threshold proportional to the number of points in the contour. The current implementation takes these two variables (one for minimum number of points in a corner, one for the distance between points) as user specified parameters and leaves any possible mathematical relationships aside for now.
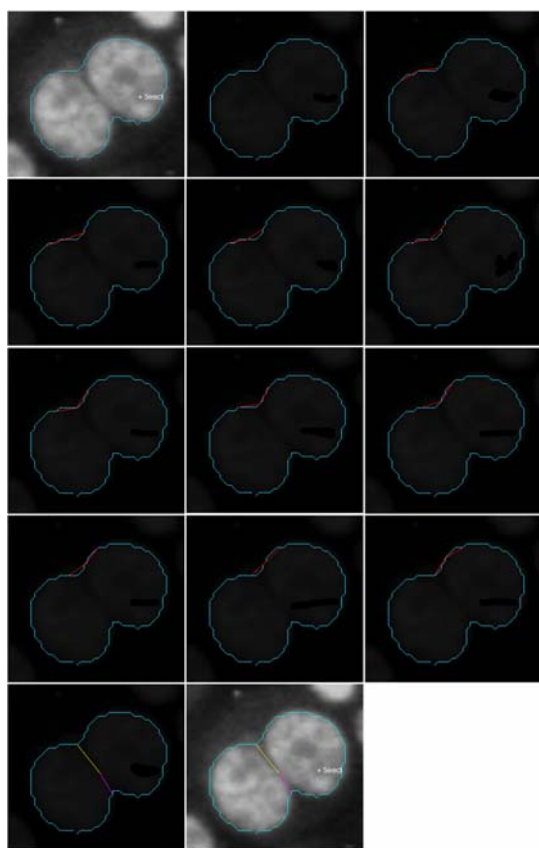
FIG 6. A CONTOUR CONSISTS OF A SERIES OF POINTS. A POINT P, IS CONNECTED TO POINTS FIVE STEPS EITHER SIDE OF ITSELF. IN SLIDE C THIS PRODUCES A CONVEX ANGLE. IN D A STRAIGHT LINE IS PRODUCED. E - J SHOW CONCAVE ANGLES, K A STRAIGHT LINE AND L PRODUCE A CONVEX ANGLE AGAIN, THUS TERMINATING THE SERIES. TAKING JUST THE CONCAVE ANGLES E-J, J IS IGNORED SO THAT THERE IS A CENTRAL POINT (G) WITH AN EVEN NUMBER OF POINTS ON EITHER SIDE OF IT

In FIG 6, we can see the effect of stepping through an actual corner with lines drawn to co-ordinates five points away from the current point. In slide J, all the points which are one step away from the current point would form a convex angle. This is true for a number of other points and so the minimum number of concave angles in series would not be reached. In fact at point G, were the centre of the corner has been established we would have a concave angle. However, F and H would not, and the number of points in the corner insufficient for it to be considered a concave region. Similarly if we were to continue to iterate across the contour, when we reached the right side of the shape we would find a number of individual concave points that would be falsely interpreted.

When extrapolating the boundary, finding the centre point of each region and connecting them resulted in a straight line and was considered too simplistic. Due to the method of acquiring angles as explained in FIG 5, the angle of points, or the mean that could be calculated from a number of angles, wasn't considered to bear enough relation to the actual contour to be useful in extrapolating a corner. In trying to establish a centre point for the line connecting these two corners, two lines were drawn; one from the start point of one corner to the start point of another, and the other from the finish points (points were iterated over in a clockwise manner). Using the method explained above the intersection of these two lines could be used to determine an extra point that could be used to extrapolate the corner (FIG 7). This idea evolved to connect each point in one corner with corresponding points in another.

The technique produced a set of points that were not aligned in a straight line and therefore considered more realistic. The number of points in each corner was balanced beforehand to enable stepping over the points easier. In cases where the number of points in a corner was low and insufficient to use the method, the middle of the line connected to the centre point of each corner was used instead. Finally, a third user variable was introduced that determined which points should be used either side of the corner centres in order to compensate for any smoothing of the corner that was introduced by the original image processing filters (visible in the leftmost image in FIG 8). The final stage simply involved producing two separate polygons each using the half the original's co-ordinates and the new extrapolated line.
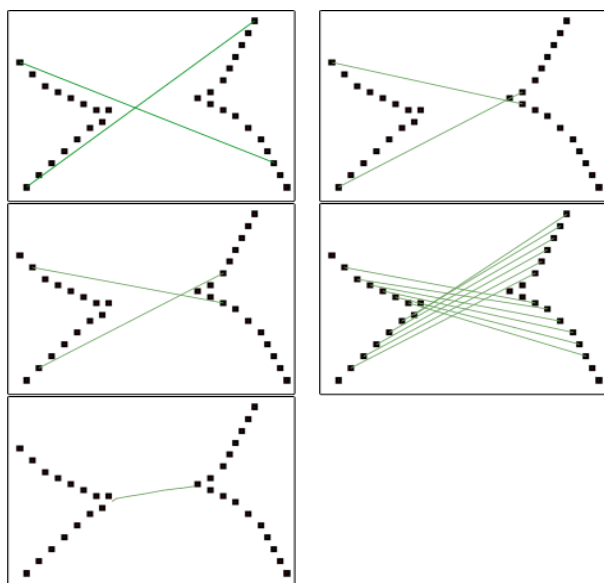
FIG 7 FROM LEFT TO RIGHT, TOP TO BOTTOM. A) ORIGINAL ATTEMPT TO CONNECT POINTS TO EXTRAPOLATE EDGE. AN ADDITIONAL POINT WAS PUT DOWN WHERE THE LINES INTERSECT AND THE CENTRE POINT OF EACH CORNER CONNECTED WITH IT. B) IN THE REVISED VERSION, THE POINT ONE STEP AWAY FROM THE CENTRE IS CONNECTED TO THE POINT FURTHEST FROM THE CENTRE ON THE OPPOSITE SIDE. THE INTERSECTION IS RECORDED. C) LINES ARE DRAWN FROM ALL THE POINTS WHICH ARE ONE STEP FURTHER AWAY FROM THE CORNER ON ONE SIDE AND ONE STEP CLOSER ON THE OTHER. D) PROCESS IS REPEATED FOR EACH POINT. E) INTERSECTIONS ARE JOINED TOGETHER TO PRODUCE EXTRAPOLATED EDGE
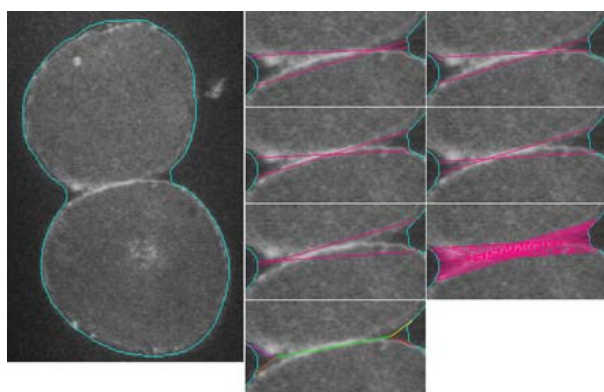
FIG 8 LEFTMOST IMAGE SHOWING TWO NEIGHBOURING CELLS AND THE FALSELY IDENTIFIED CONTOUR FROM CONFIDENCE CONNECTED THRESHOLDING. LINES ARE DRAWN FROM THE INNERMOST POINTS TO THE OUTERMOST POINTS AND THE INTERSECTION MARKED AS IN FIGURE 7. FINALLY AFTER THE LINE IS EXTRAPOLATED, IT IS CONNECTED TO POINTS A NUMBER OF STEPS EITHER SIDE OF EACH CORNERS CENTRE (NUMBER OF STEPS PROVIDED BY THE USER)

Results

Various examples of the successful implementation of the technique have been included below. From left to right in each image sequence we find the original,

unaltered image; the result of watershed segmentation; over-segmentation from thresholding, under-segmentation from thresholding; and finally the application of concave corner detection to the under-segmented threshold contour.

In each of the examples of over-segmentation, the images fall into two main categories. Either the region of one cell has expanded into the neighbouring cell before fully encompassing the intended cell (FIGs 9, 10, 12, 13, and 14) or has been captured just prior to overflowing showing the best possible result (FIG 11). Similarly, watersheds have largely been captured prior to over-segmentation at the earliest point (i.e. lowest water level) that the complete cell area has been detected (FIGs 9, 10, 11, 12, 13). It is possible that changes to the watershed level will succeed in distinguishing neighbouring cells but is often imperfect (compare for example the watershed in FIG 11 against FIG 2c; although the same two cells have been successfully distinguished in FIG 2c, the edge between them is inaccurate). FIGs 15 and 16 display cases where over-segmentation has begun before the entire cell area has been identified. FIG 14 shows a complete failure for the watershed to identify the boundary between two cells.
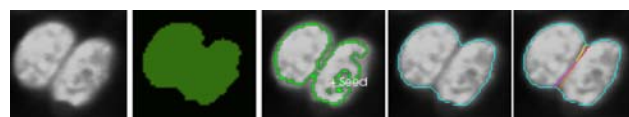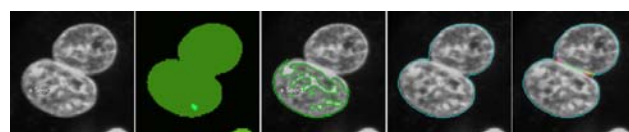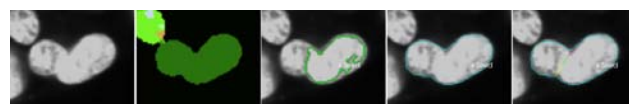

FIG 9
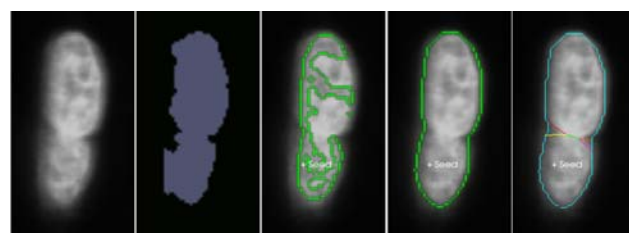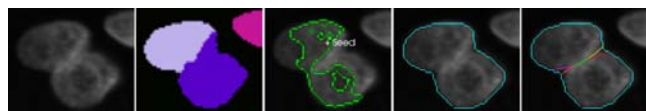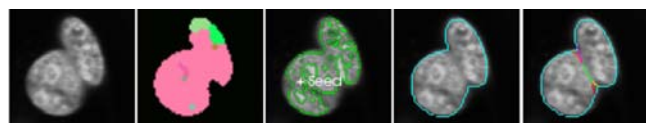

FIG 10

FIG 11

FIG 12
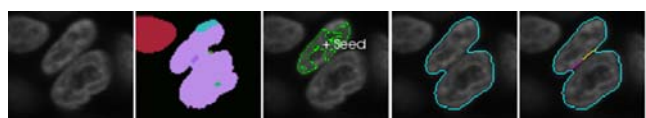

FIG 13

FIG 14

FIG 15

FIG 16

Under-segmented thresholds are largely accurate although sometimes the contour does not produce as severe an angle as is visible in the concave regions either side of the boundary (most clear in FIG 8). For this reason, the optional ability for the user to connect the extrapolated edge to points either side of the corner centre was deemed appropriate (FIGs 11, 13, 14, 15, 17). Note that FIG 11 and FIG 16 show satisfying results of the same cells seen in FIGs 2 and 3. In FIG 9 and 16, concave regions consisted of only three points and so the line extrapolation technique results in only a single point joined to either side by connecting lines.
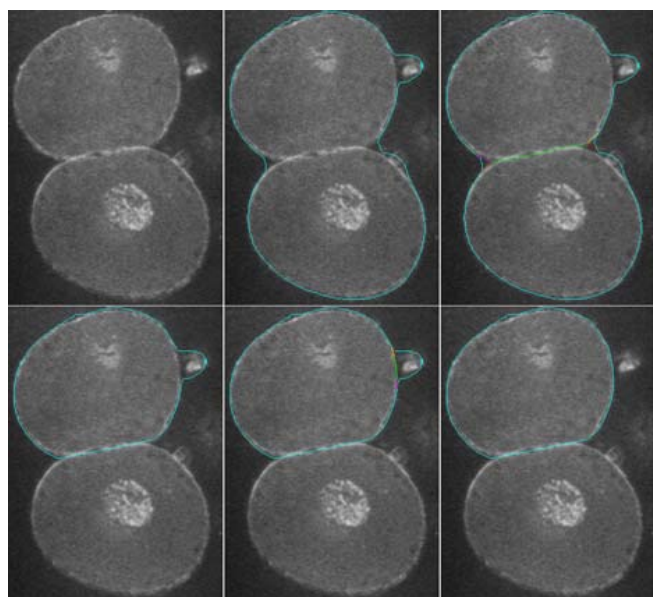
FIG 17 CAPABILITY OF SYSTEM TO PRIORITISE CORNERS BY SIZE. IN THE TOP THREE SLIDES, AN ANOMALY ON THE RIGHT OF THE UPPER CELL CREATES ADDITIONAL CONCAVE CORNERS IN THE PERIMETER WHICH ARE IGNORED IN FAVOUR OF THE LARGER REGIONS AT THE CENTRE BETWEEN THE TWO CELLS. IN THE BOTTOM THREE SLIDES, AFTER THE CELLS ARE DIVIDED, THE CORNERS OF THE ANOMALY ARE USED TO CORRECT THE TOP CELL'S CONTOUR.

Conclusion

The technique proposed above suggests a semi-automatic technique that can use just the contour of a shape to accurately estimate concave regions and identify possible boundaries. It requires user provided values to be effective but relationships between the total number of coordinates in a contour and these values seems likely, potentially allowing for automation of the procedure. The technique successfully compensates for noise and requires no prior knowledge of an objects intended shape. Although successful use of this information is restricted when more than two cells are adjacent to each other, it is hoped further investigation would likely expose methods to overcome this problem, particularly if the identified concave regions could be fed back into the original threshold filter and provide information to affect the inclusion of pixels that would compensate for the original under-segmentation.

REFERENCES

Filho E.S, Noble. J.A, Wells. D., 2010, "A Review on Automatic Ananlysis of Human Embryo Microscope Images"

Gardner D. K., Surrey E., Minjarez D., Leitz A., Stevens J., Schoolcraft W.B., (2004) "Single blastocyst transfer: a prospective randomized trial", Fertil. Steril., vol 81, no 3. pp. 551-555

Giusti A, Corani G, Gambardella L, Magli C, and Gianaroli L, "Blasotmere segmentation and 3D morphology measurements of early embryos from hoffman

modulation contrast image stacks", 2010 IEEE International Symposium on Biomedical Imaging. 13-17 April 2010, Rotterdam,The Netherlands.

Karlsson A, Overgaard NC, and Heyden A, 2004 "Automatic segmentation of zona pellucida in hmc images of human embryos", Proceedings of the 17th International Conference on Pattern Recognition (ICPR'04), pp. 23-26

Morales DA, Bengoetxea E, and Larranaga P, 2008 "Automatic segmentation of zona pellucida in human embryo images applying an active contour model" Proceedings of the 12th AnnualConference on Medical Image Understanding and Analysis, pp.209-213

Pedersen UD, Olsen OF, Olsen NH, 2003 "A multiphase variational level set approach for modelling human embryos", Proceedings of the 2nd IEEE Workshop on variational, Geometric and Level Set Methods

Scott LA, Smith S, 1998, "The successful use of pronclear embryo transfers the day following oocyte retrieval", Human Reproduction vol113, no. 4, pp. 1003-1013

Smochina C, Serban A, Manta V (2011) Segmentation of cell nuclei within chained structures in microscopic images of colon sections

Steer C.V, Mills C.L, Tan S.L, Campbell S, Edwards R.G, 1992 "The cumulative embryo score: a predictive embryo scoring technique to select the optimal number of embryos to transfer in an in-vitro fertilization and embryo transfer programme", Human Reproduction, vol. 7, pp. 117-119

Tassy, O., Daian, F., Hudson, C., Bertrand, V. and Lemaire, P. (2006) "A Quantitative Approach to the Study of Cell Shapes and Interactions During Early Chordate Embryogenesis." Current Biology, 16 pp. 345-358.

Kieran Rafferty is a MSc student at Manchester Metropolitan University, Manchester, UK.

Sarah Drury is a final year PhD student at Warwick Medical School, University of Warwick (Division of Reproductive Health), and a Clinical Embryologist at the Centre for Reproductive Medicine, University Hospitals Coventry and Warwickshire. Her research is focussed on human preimplantation embryo development and viability.

Geraldine Hartshorne is Professorial Fellow at Warwick Medical School, University of Warwick, UK, and Scientific Director of the Centre for Reproductive Medicine, University Hospitals Coventry and Warwickshire NHS Trust. Her research interests include human and embryo development, with a view to predicting pregnancy. She has established Royal College examinations for clinical embryologists and is keen on developing technological simulations for embryology training, education and examination.

Silvester Czanner is a lecturer in Computer Games at Manchester Metropolitan University in UK. He is doing research in areas of Computer Graphics, Medical Imaging and Digital Technologies, especially their applications in education such as using graphics, images and animations to enable and enhance learning. Dr. Czanner is a senior member of the Association for Computing Machinery (ACM).